

CLAIMS

We claim:

Sub
A1

1. A method for maintaining cache coherence in a multiprocessor system having a plurality of nodes coupled by an interconnecting communications pathway which is capable of storing information regarding the location and state of data within the system, each node having at least one cache, a memory device local to the node, and at least one processor device, the processor device within each node being capable of accessing data from the local memory device, the local cache, or over the interconnecting communications pathway from a non-local memory device or a non-local cache, the method comprising:

storing information regarding the state of data in said interconnecting pathway;
checking said stored information to determine the location of the most current copy of a requested portion of data, in response to a request by a requesting node for the requested portion of data;

retrieving said current copy of requested portion of data and directing said data to the requesting node;

checking said stored information to determine the location of the requested data; and
directing the system to send said requested data to the requesting node without going through the said interconnecting communications pathway.

2. A multiprocessor computer system comprising:
a plurality of nodes, each node including at least one processor and portion of a shared distributed system memory coupled to said processor; and
a communication pathway connecting said nodes and including a central hardware device which stores location and state information of data stored in the memory of the nodes.

3. The multiprocessor system of claim 2, wherein each node includes memory accessible to it without communications through said communications pathway, and memory accessible remotely by others of the nodes.

4. The multiprocessor system of claim 2 wherein further said central hardware device stores information for determining which nodes or processors are storing copies of one or more identified data in each said node's memory.

5. The multiprocessor system of claim 2 wherein said central hardware device compares requested data with the stored location and the state of data in the nodes, directs requested data to the requesting node, and sends requests for additional data to other nodes for which said device stores the location of data.

6. The multiprocessor system of claim 5 wherein said central hardware device includes a dispatch buffer operatively connected to the nodes, and issues requests for information related to the state of identified data to other nodes simultaneously with the communication of data to a target node.

7. The multiprocessor system of claim 5 wherein said interconnecting communications pathway includes a first pathway storing the location and state of data in the nodes, and a second pathway communicating the data requested by said target node.

8. A method for maintaining cache coherence in a multiprocessor system having a plurality of nodes coupled by an interconnecting communications pathway which is capable of storing information regarding the location and state of data within the system, each node having at least one cache, a memory device local to the node and at least one processor device, the

memory and processor device being coupled to form a complete subsystem, the processor device within each node being capable of accessing data from the local memory device, the local cache, or over the interconnecting communications pathway from a non-local memory device, or a non-local cache, wherein further said communications pathway is comprised of a first pathway which communicates the state and location of data within the nodes, and a second communications pathway which communicates the data between the nodes, the method comprising:

storing information regarding the state of data in said first pathway; said first pathway checking said stored information to determine the location of the most current copy of a requested portion of data, in response to a request by a requesting node for data;

said first pathway directing said second pathway to forward the said most current copy of said data to the requesting node; and

said second pathway retrieving said current copy of requested portion of data and directing said data to a target node.

9. The method of claim 8 wherein said step of said first pathway checking said stored information to determine the location of the most current copy of a requested portion of data, in response to a request by a requesting node for data, comprises:

storing information about the state of data in each node in said first communications pathway;

checking the state of requested data stored in each node upon request for the data from a node by reading the said stored information and determining the desired state defined as the most current copy of said stored data.